

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Systems And Methods For Replicating Virtual Memory
On A Host Computer And Debugging Using Replicated
Memory**

Inventors:

Gregory Hogdal
John Eldridge

ATTORNEY'S DOCKET NO. MS1-671US

RELATED APPLICATIONS

The present application is a continuation of U.S. Patent Application Serial Number 09/865,934 filed 5/24/2001, by the Applicants named herein and entitled "Systems and Methods For Replicating Virtual Memory On A Host Computer And Debugging Using Replicated Memory." The present application is also related to U.S. Provisional Patent Application Serial Number 60/234,643 filed 9/22/2000 by the Applicants named herein and entitled "Systems and Methods For Replicating Virtual Memory On A Host Computer And Debugging Using Replicated Memory."

TECHNICAL FIELD

The systems and methods described herein relate to debugging computing systems and, more particularly, to debugging a target computing system replicating target virtual memory translation on a host computer and debugging on the host.

BACKGROUND

As computing technology has advanced, the size of computer software applications and the operating systems that run them has grown larger and larger. As the size of the software has increased, so had the demand that is placed on the memory required to support these programs. To deal with the need for an increased number of addressed memory space, the concept of virtual memory was developed. Today, virtually all modern operating systems provide a form of virtual memory to applications.

1 One newly developed operating system that utilizes virtual memory
2 management is the WINDOWS CE operating system produced by MICROSOFT
3 CORP. WINDOWS CE is a lightweight operating system that is ideal for use in
4 PDAs (Personal Digital Assistants), hand-held computers, palm computers,
5 electronic appliances and the like. WINDOWS CE provides a page-based virtual
6 memory management scheme that allows applications to realize a 32-bit linear
7 address space for four (4) gigabytes (GB) of memory.

8 A system that utilizes a virtual memory scheme poses a particular problem
9 when the system fails and an attempt is made to debug the system. To debug a
10 system, a software program is used to examine the contents of the system's
11 memory and registers to determine a problem with a system. Debuggers require
12 that a virtual memory-based system that is being debugged be operational because
13 the debugging software at least requires the CPU to execute the software.
14 However, there are situations in which the system is not operational, *i.e.*, the CPU
15 will not execute the debugging code, where it is necessary or desirable to execute
16 debugging code to determine the cause of a system fault. Such a situation arises
17 when using hardware-assisted debugging equipment that completely freezes the
18 system in order to debug it, or when a snapshot of the system is taken to be
19 debugged at a later time (this is referred to as "post-mortem" debugging).

20 When such a situation arises, it is impossible to rely on the kernel of the
21 operating system to handle a page fault exception to load a missing page when a
22 virtual memory location not currently loaded is accessed for debugging purposes.
23
24
25

SUMMARY

Systems and methods are described herein that provide a means for a host computer to describe a translation equivalent of one typically performed by CPU table look-aside buffer (TLB) registers of a target computer, after the kernel of the target computer has added the mapping of a page required by the debugger, in the CPU TLB. A host-side application locally replicates the mechanism normally used by the kernel of the target computer to map a virtual address into a physical address. After such a translation is accomplished on the host computer, debugging can be performed on the host computer.

The host-side virtual to physical address translation assumes the following environment:

- (1) The target system is running on a CPU that supports fixed paged memory management;
- (2) The target system is running an operating system that enables and uses the paged memory management;
- (3) The target operating system memory management is table driven or has hard-coded logic; and
- (4) The tables used by the target operating system memory management (if applicable) are located either in a known address range of physical memory, or in a known address range of the virtual memory whose pages have been locked (so their addresses are translatable by looking up the current CPU TLB).

In accordance with the present invention, the host-side virtual to physical address translation performs the following major tasks:

- (1) determining if the memory management of the CPU is enabled or not
(in which case the following tasks are not necessary);
- (2) locating the data (tables) used by the kernel of the target system
directly to perform the translation;
- (3) replicating locally the data (tables) used by the kernel of the target
system to perform the translation;
- (4) checking the validity of the data (tables) used by the kernel of the target
system to perform the translation;
- (5) in the case that the data (tables) used by the kernel of the target system
to perform the translation are valid, they can be used to perform the
translation on the host side; and
- (6) cache the data for future use (optional).

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of exemplary methods and arrangements of the present invention may be had by reference to the following detailed description when taken in conjunction with the accompanying drawings wherein:

Fig. 1 is an exemplary computer system on which the present invention may be implemented.

Fig. 2 is a block diagram of a host computer and a target computer, the diagram illustrating one implementation of the invention.

Fig. 3 is a block diagram of a host computer and a target computer.

Fig. 4 is a flow diagram depicting a method for replicating and translating virtual address data from a target computer on a host computer.

DETAILED DESCRIPTION

The invention is illustrated in the drawings as being implemented in a suitable computing environment. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, to be executed by a computing device, such as a personal computer or a hand-held computer or electronic device. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

Exemplary Computer Environment

The various components and functionality described herein are implemented with a number of individual computers. Fig. 1 shows components of typical example of such a computer, referred by to reference numeral 100. The components shown in Fig. 1 are only examples, and are not intended to suggest any limitation as to the scope of the functionality of the invention; the invention is not necessarily dependent on the features shown in Fig. 1. In addition, the fact that a personal computer and its components is depicted in Fig. 1 is exemplary

1 only and is not meant to limit the scope of the invention or inventions described
2 herein. For example, one or more implementations described herein may utilize a
3 handheld computer. Those skilled in the art will appreciate the environment
4 required to implement the systems and methods described herein.

5 Generally, various different general purpose or special purpose computing
6 system configurations can be used. Examples of well known computing systems,
7 environments, and/or configurations that may be suitable for use with the
8 invention include, but are not limited to, personal computers, server computers,
9 hand-held or laptop devices, multiprocessor systems, microprocessor-based
10 systems, set top boxes, programmable consumer electronics, network PCs,
11 minicomputers, mainframe computers, distributed computing environments that
12 include any of the above systems or devices, and the like.

13 The functionality of the computers is embodied in many cases by
14 computer-executable instructions, such as program modules, that are executed by
15 the computers. Generally, program modules include routines, programs, objects,
16 components, data structures, etc. that perform particular tasks or implement
17 particular abstract data types. Tasks might also be performed by remote
18 processing devices that are linked through a communications network. In a
19 distributed computing environment, program modules may be located in both local
20 and remote computer storage media.

21 The instructions and/or program modules are stored at different times in the
22 various computer-readable media that are either part of the computer or that can be
23 read by the computer. Programs are typically distributed, for example, on floppy
24 disks, CD-ROMs, DVD, or some form of communication media such as a
25 modulated signal. From there, they are installed or loaded into the secondary

1 memory of a computer. At execution, they are loaded at least partially into the
2 computer's primary electronic memory. The invention described herein includes
3 these and other various types of computer-readable media when such media
4 contain instructions programs, and/or modules for implementing the steps
5 described below in conjunction with a microprocessor or other data processors.
6 The invention also includes the computer itself when programmed according to
7 the methods and techniques described below.

8 For purposes of illustration, programs and other executable program
9 components such as the operating system are illustrated herein as discrete blocks,
10 although it is recognized that such programs and components reside at various
11 times in different storage components of the computer, and are executed by the
12 data processor(s) of the computer.

13 With reference to Fig. 1, the components of computer 100 may include, but
14 are not limited to, a processing unit 120, a system memory 130, and a system bus
15 121 that couples various system components including the system memory to the
16 processing unit 120. The system bus 121 may be any of several types of bus
17 structures including a memory bus or memory controller, a peripheral bus, and a
18 local bus using any of a variety of bus architectures. By way of example, and not
19 limitation, such architectures include Industry Standard Architecture (ISA) bus,
20 Micro Channel Architecture (MCA) bus, Enhanced ISA (EISAA) bus, Video
21 Electronics Standards Association (VESA) local bus, and Peripheral Component
22 Interconnect (PCI) bus also known as the Mezzanine bus.

23 Computer 100 typically includes a variety of computer-readable media.
24 Computer-readable media can be any available media that can be accessed by
25 computer 100 and includes both volatile and nonvolatile media, removable and

1 non-removable media. By way of example, and not limitation, computer-readable
2 media may comprise computer storage media and communication media.
3 "Computer storage media" includes both volatile and nonvolatile, removable and
4 non-removable media implemented in any method or technology for storage of
5 information such as computer-readable instructions, data structures, program
6 modules, or other data. Computer storage media includes, but is not limited to,
7 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
8 digital versatile disks (DVD) or other optical disk storage, magnetic cassettes,
9 magnetic tape, magnetic disk storage or other magnetic storage devices, or any
10 other medium which can be used to store the desired information and which can be
11 accessed by computer 110. Communication media typically embodies computer-
12 readable instructions, data structures, program modules or other data in a
13 modulated data signal such as a carrier wave or other transport mechanism and
14 includes any information delivery media. The term "modulated data signal"
15 means a signal that has one or more of its characteristics set or changed in such a
16 manner as to encode information in the signal. By way of example, and not
17 limitation, communication media includes wired media such as a wired network or
18 direct-wired connection and wireless media such as acoustic, RF, infrared and
19 other wireless media. Combinations of any of the above should also be included
20 within the scope of computer readable media.

21 The system memory 130 includes computer storage media in the form of
22 volatile and/or nonvolatile memory such as read only memory (ROM) 131 and
23 random access memory (RAM) 132. A basic input/output system 133 (BIOS),
24 containing the basic routines that help to transfer information between elements
25 within computer 100, such as during start-up, is typically stored in ROM 131.

1 RAM 132 typically contains data and/or program modules that are immediately
2 accessible to and/or presently being operated on by processing unit 120. By way
3 of example, and not limitation, Fig. 1 illustrates operating system 134, application
4 programs 135, other program modules 136, and program data 137.

5 The computer 100 may also include other removable/non-removable,
6 volatile/nonvolatile computer storage media. By way of example only, Fig. 1
7 illustrates a hard disk drive 141 that reads from or writes to non-removable,
8 nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to
9 a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that
10 reads from or writes to a removable, nonvolatile optical disk 156 such as a CD
11 ROM or other optical media. Other removable/non-removable,
12 volatile/nonvolatile computer storage media that can be used in the exemplary
13 operating environment include, but are not limited to, magnetic tape cassettes,
14 flash memory cards, digital versatile disks, digital video tape, solid state RAM,
15 solid state ROM, and the like. The hard disk drive 141 is typically connected to
16 the system bus 121 through an non-removable memory interface such as interface
17 140, and magnetic disk drive 151 and optical disk drive 155 are typically
18 connected to the system bus 121 by a removable memory interface such as
19 interface 150.

20 The drives and their associated computer storage media discussed above
21 and illustrated in Fig. 1 provide storage of computer-readable instructions, data
22 structures, program modules, and other data for computer 100. In Fig. 1, for
23 example, hard disk drive 141 is illustrated as storing operating system 144,
24 application programs 145, other program modules 146, and program data 147.
25 Note that these components can either be the same as or different from operating

1 system 134, application programs 135, other program modules 136, and program
2 data 137. Operating system 144, application programs 145, other program
3 modules 146, and program data 147 are given different numbers here to illustrate
4 that, at a minimum, they are different copies. A user may enter commands and
5 information into the computer 100 through input devices such as a keyboard 162
6 and pointing device 161, commonly referred to as a mouse, trackball, or touch
7 pad. Other input devices (not shown) may include a microphone, joystick, game
8 pad, satellite dish, scanner, or the like. These and other input devices are often
9 connected to the processing unit 120 through a user input interface 160 that is
10 coupled to the system bus, but may be connected by other interface and bus
11 structures, such as a parallel port, game port, or a universal serial bus (USB). A
12 monitor 191 or other type of display device is also connected to the system bus
13 121 via an interface, such as a video interface 190. In addition to the monitor,
14 computers may also include other peripheral output devices such as speakers 197
15 and printer 196, which may be connected through an output peripheral interface
16 195.

17 The computer may operate in a networked environment using logical
18 connections to one or more remote computers, such as a remote computer 180.
19 The remote computer 180 may be a personal computer, a server, a router, a
20 network PC, a peer device or other common network node, and typically includes
21 many or all of the elements described above relative to computer 100, although
22 only a memory storage device 181 has been illustrated in Fig. 1. The logical
23 connections depicted in Fig. 1 include a local area network (LAN) 171 and a wide
24 area network (WAN) 173, but may also include other networks. Such networking
25

1 environments are commonplace in offices, enterprise-wide computer networks,
2 intranets, and the Internet.

3 When used in a LAN networking environment, the computer 100 is
4 connected to the LAN 171 through a network interface or adapter 170. When used
5 in a WAN networking environment, the computer 100 typically includes a modem
6 172 or other means for establishing communications over the WAN 173, such as
7 the Internet. The modem 172, which may be internal or external, may be
8 connected to the system bus 121 via the user input interface 160, or other
9 appropriate mechanism. In a networked environment, program modules depicted
10 relative to the computer 100, or portions thereof, may be stored in the remote
11 memory storage device. By way of example, and not limitation, Fig. 1 illustrates
12 remote application programs 185 as residing on memory device 181. It will be
13 appreciated that the network connections shown are exemplary and other means of
14 establishing a communications link between the computers may be used.

15 WINDOWS NT Virtual Memory System

16 The WINDOWS NT virtual memory system has been available for several
17 years and is well known in the art. WINDOWS CE utilizes a similar virtual
18 memory system. A program running on WINDOWS NT or WINDOWS CE can
19 utilize 32 bits of address space. All programs running on WINDOWS CE share a
20 common 32-bit address space. This translates to four (4) gigabytes (GB) of virtual
21 memory. The upper half of this is devoted to system code and data and is only
22 visible to the process when it is in privileged mode. The lower half (2 GB) is
23 available to the user program when it is in user mode, and to those user-mode
24 system services called by the program. On WINDOWS CE, each application gets
25 a 32 MB slot from the lower 2 GB of user space.

1 The RAM (Random Access Memory) of a computer running WINDOWS
2 NT or CE is divided into two categories: non-paged and paged. Non-paged code
3 or data must stay in memory and cannot be written to or retrieved from
4 peripherals. Peripheral include disks, a local area network (LAN), a CD-ROM,
5 and other devices. Paged memory is RAM which the system can use and later
6 reuse to hold various pages of memory from peripherals. Paged memory is
7 divided into page frames, that hold various pages from time to time.

8 Page size varies with the computer's processor type. For example, page
9 size is 4096 bytes (4K) for 386, 486 and Pentium-class processors, and the same
10 for MIPS and ARM processors. When a page of code or data is required from a
11 peripheral, the memory manager finds a free page frame in which to place the
12 required page. The system transfers the required page, and processing continues.
13 If no page frame is free, the memory manager must select one to reuse. The
14 memory manager tries to find a page frame whose contents have not been
15 referenced for a while. When the memory manager finds a suitable page frame, it
16 discards the page in it.

17 Normally, programs execute by fetching one instruction after another from
18 a code page (a page that contains program instructions) until they call or return to
19 a routine in some other code page or make a jump to code in another page. Or,
20 they can simply run off the end of the current page and need the next page. Such a
21 transfer of instruction control to a new page causes a page fault if the needed page
22 is not currently in the working set of the process. The working set of the process
23 is the set of pages currently visible to the process in RAM.

24 A page fault can be resolved quickly if the memory manager finds the page
25 elsewhere in RAM. It might be in the working set of some other process or

1 processes, or it might have been removed from the current process's working set
2 by the memory manager in an overzealous attempt to keep the process trim and fit.
3 The memory manager places such pages on a list of page frames called the
4 standby list, and they can be reinserted into the working set of a process. But if
5 the page is not in RAM, the memory manager must find a free page frame, or
6 make one free as described above, and then fetch the required page from the
7 peripheral. One characteristic of code pages is that it isn't normal for code to be
8 modified while in RAM, so code pages can be discarded without being written
9 back to disk.

10 Data pages, which contain data used by a program, are accessed in a
11 somewhat more random fashion than code pages. Each instruction in a program
12 can reference data allocated anywhere in the address space of a process. The
13 principle, however, is much the same. If an attempt is made to access a data page
14 not in the working set of the process, a page fault occurs. From that point on, the
15 process is just as described for code pages.

16 Preferred Implementation

17 Fig. 2 is a high-level block diagram of a system in accordance with one
18 implementation of the present invention, which will be used to discuss a broad
19 overview of the invention. A host computer 200 includes memory 202. A
20 debugger 204 is stored in the memory 202 of the host computer 200. A target
21 computer 206 includes memory 208 that stores several translation tables 210.

22 The host computer 200 accesses the target computer 206 via an access
23 mechanism 212 such as hardware-assisted debug probes. The translation tables
24 210 of the target computer 206 are replicated (translation tables 210') in the
25 memory 202 of the host computer 200 by performing a similar operation that is

1 performed by the target computer to realize the virtual memory. The translation
2 tables 210' can then be used to translate physical memory addresses to virtual
3 memory addresses. In one implementation, the virtual memory addresses are
4 stored in the memory 202 of the host computer 200 after they are translated. The
5 debugger 204 is now able to read the physical memory corresponding to the
6 virtual addresses (translating) it required to access to potentially determine a cause
7 of a fault in the target computer, by analyzing data produced from the replication.

8 Fig. 3 is a more detailed block diagram of a host computer 300 and a target
9 computer 302 as implemented in the current invention. The host computer 300
10 includes a processor 304, memory 306 and cache memory 308. The memory 306
11 of the host computer 300 stores an operating system 310 that executes on the
12 processor 304, a debugger 312, and an address table 313. The memory 306 also
13 includes a data retrieval component 314, an address translation component 316,
14 and a memory management identifier 318. The function of these components will
15 be described in greater detail below.

16 The target computer 302 includes a processor 320 having a register 322.
17 The register 322 may be one of several registers in the processor 320. The
18 processor 320 of the target computer 302 supports fixed paged memory
19 management. The target computer 302 also includes memory 324 that stores an
20 operating system 326 and virtual address data 328. The operating system 326 uses
21 the paged memory management that is supported by the processor 320. In the
22 described implementation example, the operating system 326 is table driven,
23 although in other implementations, the operating system may have hard-coded
24 logic. A data link 330 enables data transfer between the host computer 300 and the
25 target computer 302.

1 Although the systems and methods described herein can be implemented in
2 numerous systems that utilize virtual memory management, the present discussion
3 will use many terms and functions specific to the virtual memory management
4 system that is utilized in the WINDOWS CE and WINDOWS NT operating
5 systems produced by MICROSOFT CORPORATION. However, this is not
6 intended to limit the scope of the invention to these specific products.

7 Fig. 4 is a flow diagram of a method to replicate virtual memory data from
8 the target computer 302 on the host computer 300. Continuing reference will be
9 made to the elements and reference numerals contained in Fig. 3. The following
10 discussion assumes that a connection has been made between the host computer
11 300 and the target computer 302.

12 At step 400, the memory management identifier determines if the memory
13 management unit of the target processor 320 is enabled. This process is specific to
14 the type of processor 320 that is in the target compute 302. Typically, making this
15 determination requires locating and reading a value contained in the register 322
16 of the processor 320. If the memory management unit is disabled ("No" branch,
17 step 400), then the procedure cannot continue.

18 If memory management is enabled on the processor 320 ("Yes" branch, step
19 400), then the data retrieval component 314 locates the data used by the target
20 computer 302 to perform memory translations (step 402). This requires knowing
21 exactly where in the memory 324 the virtual address data 328 is stored. If only the
22 virtual address is known, then the translation can be performed on the host
23 computer 300 by looking at content of a certain register 322 in the target processor
24 (CPU Table Look-aside Buffer (TLB) register). The absolute location is normally
25

1 given by accessing a binary file that contains the image downloaded on the target
2 computer 302 and certain symbolic debug information.

3 At step 404, the address translation component 316 then reads the virtual
4 address data 328 on the target computer 302 and replicates the data on the host
5 side at step 406. This is the data that is used by the target computer 302 to
6 perform virtual memory translations. This procedure is done through an available
7 target access application program interface (API) such as the eXDI API in
8 Platform Builder for WINDOWS CE 3.0, that typically makes use of hardware-
9 assisted equipment to read the target memory 324 directly. The virtual address
10 data 328 typically comprises tables that store data that can be used to determine a
11 virtual address that is represented by a physical address. However, depending on
12 the type of virtual memory management utilized on the target computer, the virtual
13 address data 328 may comprise some other sort of data.

14 In the preferred implementation, the virtual address data 328 is validated at
15 step 408. This is necessary to prevent false translation in the case that the target
16 operating system 366 is not initialized, is partially initialized, corrupted, or in the
17 process of updating its tables. This can be achieved by any method known in the
18 art, such as by using redundant information and markers within fields of tables. If
19 the virtual address data 328 is not valid, then the processing is termination ("No"
20 branch, step 408). If the virtual address data 328 is valid ("Yes" branch, step 410),
21 then the processing continues at step 412.

22 At step 412, the virtual address data 328 is translated by the address
23 translation component 316 of the host computer 300. This process takes the
24 contents of the physical memory 324 of the target computer 302 and applies the
25 same process that the target computer 302 operating system 326 uses to utilize

1 virtual memory addressing. The details of this task are specific to the virtual
2 memory management scheme utilized by the operating system 326 and each
3 particular VMM scheme will be appreciated by those skilled in the art.

4 Also at step 412, the debugger 312 is used to debug a fault on the target
5 computer 302 by using the translated virtual memory data. The presence of a fault
6 to debug is not required for the present invention, but is indicated as the typical
7 debugger usage step. This step is shown as being combined with the translation
8 because the debugging can be done as the virtual memory data 328 is being
9 translated. In one implementation, the translated virtual memory data is cached in
10 the cache memory 308 at step 420. This can be done either before or after the
11 debugging process. If done before, the data is cached and the debugging is
12 performed on the data stored in the cache memory 308. If done after, the
13 debugging is done as the data is being translated, then the translated data is cached
14 so it can be referenced at a later time, if necessary.

15 Although not required, caching is desirable because it speeds up success
16 translations. Some data can be saved permanently for an execution session (after
17 initialization of the target operating system 326 to the next reset), such as page size
18 and the root to the virtual address data 328. All the other data can be cached while
19 the processor 320 on the target computer 302 is halted. The replication of linked
20 data structure (tables containing pointers to other tables) implies "fixing up" the
21 pointers as the pointed tabled are replicated.

Conclusion

The system and methods as described, thus provide a manner in which a halted system can be debugged by replicating the virtual memory data from the target system on the host system. After the virtual memory data is replicated on the host system, debugging is performed on the data on the host system. This simulates debugging on the target computer, which cannot be performed directly because the processor on the target computer has halted execution.

Although details of specific implementations and embodiments are described above, such details are intended to satisfy statutory disclosure obligations rather than to limit the scope of the following claims. Thus, the invention as defined by the claims is not limited to the specific features described above. Rather, the invention is claimed in any of its forms or modifications that fall within the proper scope of the appended claims, appropriately interpreted in accordance with the doctrine of equivalents.